

Supervised Autonomy for Communication-degraded Subterranean Exploration by a Robot Team

Kyohei Otsu, Scott Tepsuporn, Rohan Thakker, Tiago Stegun Vaquero, Jeffrey A. Edlund, William Walsh, Gregory Miles, Tristan Heywood, Michael T. Wolf, and Ali-akbar Agha-mohammadi
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA 91109,
{firstname.lastname}@jpl.nasa.gov

Abstract—The importance of autonomy in robotics is magnified when the robots need to be deployed and operated in areas that are too dangerous or not accessible for humans, ranging from disaster areas (to assist in emergency situations) to Mars exploration (to uncover the mystery of our neighboring planet). The DARPA Subterranean (SubT) Challenge presents a great opportunity and a formidable robotics challenge to foster such technological advancement for operations in extreme and underground environments. Robot teams are expected to rapidly map, navigate, and search underground environments including natural cave networks, tunnel systems, and urban underground infrastructure. Subterranean environments pose significant challenges for manned and unmanned operations due to limited situational awareness. In the first phase of the DARPA Subterranean Challenge (held in August 2019; targeting underground tunnels and mines), Team CoSTAR, led by NASA JPL, placed second among 11 teams across the world, accurately mapping several kilometers of two mine systems and localizing 17 target objects in the course of four one-hour missions.

While the main goal of Team CoSTAR at the end of this three-year challenge (August 2021) is a fully autonomous robotic solution, this paper describes Team CoSTAR’s results in the first phase of the challenge (August 2019), focusing on supervised autonomy of a multi-robot team under severe communication constraints. This paper also presents the design and initial results obtained from field test campaigns conducted in various tunnel-like environments, leading to the competition.

TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. CONCEPT OF OPERATIONS.....	2
3. TECHNICAL APPROACHES	2
4. RESULTS	6
5. CONCLUSIONS.....	7
ACKNOWLEDGMENTS	8
REFERENCES	8
BIOGRAPHY	8

1. INTRODUCTION

Space agencies have long considered robotic exploration of caves, lava tubes, and subsurface voids on the Moon and Mars for scientific discovery and human precursor missions [1], [2]. Shielded from the radiation and extreme temperature variations at the surface, the natural protection in these extraterrestrial subterranean environments offers advantages for potential human habitats [3] and better preserves the history of planetary conditions and possible life [4], [5]. Further,



Figure 1. Husky robot exploring the Arch Coal mining complex in Beckley, WV

mapping their topology and topography can provide evidence about the past geological processes of the planetary bodies.

Terrestrial underground robotic exploration has gained increased attention due to the recently-initiated DARPA Subterranean (SubT) Challenge program [6]. The SubT Challenge includes a set of robotic competition events, where robot teams are expected to rapidly map, navigate, and search underground man-made and natural environments including tunnel systems, urban underground infrastructure, and natural cave networks (see representative environment in Figure 1). Environments may be as long as 8 km in length and must be explored within mission time constraints (typically 60 minutes). Only one single human supervisor (located on the surface) is allowed to supervise the entire robotic system, and communicate high-level commands to the robotic agents in this highly communication-degraded environment. Hence, autonomy is critical to the success of team.

In this paper, we present the results of Team CoSTAR (Collaborative SubTerraanean Autonomous Robots) [7] in the first year of this three year challenge. While targeting the *fully autonomous* solution at the end of the third year, in the first year’s competition (referred to as Tunnel Circuit) held in August 2019, Team CoSTAR relies on a supervised autonomy of a multi-robot team to enable rapid exploration of unknown environments. To discuss Team CoSTAR’s initial implementation, we start with several key design questions:

- **Single-robot autonomy:** What type and level of capabilities are needed for on-board autonomy?
- **Communication maintenance:** How do we maintain communication links to the robots to send commands and collect data?
- **Data prioritization:** Under the severe communication constraints, what data need to be shared between the robots and the base station to perform effective decision making?

- **Robot tasking:** How do we assign tasks to robots which might have different mobility and perception capabilities?
- **Autonomy vs human intervention:** In what level does the human supervisor intervene with the autonomous operations of the robot team?

Robotic exploration of subterranean environments poses significant challenges for operations, largely due to communications constraints. Communication in a subsurface environment has a high level of uncertainty in the reliability, capacity, and availability of the links between nodes in a network [8]. This limits a human’s situational awareness as well as the ability to send commands. While a robot may explore autonomously out of communication range for some time, a solution must balance providing the human supervisor valuable data for high-level decision-making with making efficient progress in exploring frontiers. Multi-agent systems provide particular advantages in this scenario, enabling data gathering and mapping operations into large spatial areas while exploiting a subset of robots for data relay [9], [10], [11].

Team CoSTAR operated a team of four ground robots supervised by a single human operator at the base station. Our single-robot autonomy module consists of a specification-based mission executive, large-space graph-based path planning using an information roadmap (IRM) representation, and capability-agnostic behavior engines. These capable individual robots are connected via a wireless communication backbone, built by the support robots while accounting for the geometry and signal attenuation of the environment. The human supervisor at the base station monitors the IRM, 3D volumetric map, communication graph, and location/health of all agents in a graphical user interface. The supervisor occasionally sends back a modified IRM graph to incorporate the human intention in robot’s global planning. The presented method was successfully demonstrated in simulation, indoor testing facilities, and various mines including the Tunnel Circuit. The robot team autonomously explored km-long areas in the event, accurately locating 17 objects (with less than 5 m error in location). Team CoSTAR took second place out of 11 teams across the world.

The rest of paper is organized as follows. Section 2 presents our concept of operations. Section 3 presents technical approaches. Section 4 shows the results from simulation and various field tests including the Tunnel Circuit. Finally, Section 5 concludes the paper.

2. CONCEPT OF OPERATIONS

We begin by describing our concept of operations. Each robot is assigned a role. The following roles were defined for the Tunnel Circuit: 1) Vanguard, who is responsible for exploring new regions and create a map, 2) Support, who follows the Vanguard robot to relay data, as well as building a static communication backbone to maintain links, 3) Task Allocator, who assigns missions and/or tasks to each robot. A role can be reassigned based on the mission progression. For this first Tunnel Circuit, the Task Allocator role was executed by the human supervisor.

We use the mesh networking technology to connect robots and the base station. The mesh network is formed by the base station node, robot nodes, and droppable communication nodes which are carried and deployed by carrier robots (see Figure 1). A part of mesh network is maintained with a

long communication tether that can provide stable connection regardless of environmental condition. Each node can dynamically join and leave the network.

The procedure of nominal operations is described as follows:

1. The Task Allocator sends a Vanguard robot to the environment. The Vanguard autonomously explores unknown regions, creates a map, and sends back the map to the base station.
2. The Task Allocator sends a Support robot to follow the Vanguard. The Support relays data from the Vanguard.
3. The Support robot drops a communication node if necessary. Typical drop location is a junction. It also drops a node if the signal-to-noise ratio (SNR) to the base station goes below a threshold due to distance or sharp turns.
4. Continue with the above steps until the Support deploys all available communication nodes. Then, Support becomes a new Vanguard.
5. Vanguard occasionally returns to communications by traversing the known path derived from the IRM. After regaining communications and transferring all data collected, the robot continues exploration.

A typical team configuration is a Vanguard robot and one or more Support robot(s). If three or more robots are available, different regions are assigned to each team unit to explore in parallel.

The human supervisor acts as a component in the multi-agent system. Especially, the supervisor aids the system by providing complex perception ability (e.g., detecting promising direction to explore), and making high-level strategic decisions (e.g., mission start/termination). The supervisor also engages in the recovering behavior, such as correcting global localization error or recovering immobilized vehicles.

3. TECHNICAL APPROACHES

In this section, we describe the technologies that support our concept of operations.

System Overview

The system architecture is shown in Figure 2. Our perception subsystem employs hierarchical and multi-modal approaches, including heterogeneous local state estimation [12], global localization based on pose graph optimization [13], and multi-modal artifact detection and localization [14]. The autonomy subsystem is composed of mission executive [15], capability-agnostic behavior engines, and platform-dependent mobility services which expose the common interface to the mobility behaviors. The roadmap manager takes care of maintaining sparse but information-rich representation of the environment on top of the pose graph constructed by the global localization module. Internally, these data are shared between modules using a world model concept, which enables asynchronous read/write operations from different modules. The networking subsystem is responsible to send data to inter-agent networks, which takes care of routing, data compression, and quality of service (QoS) control including reliability and timeliness.

We use ROS [16] for all robot software. To support the multi-agent operations, we set up a ROS multi-master system with open-source multi-master FKIE (MM-FKIE) package [17] and a custom cross-master messaging mechanism named multi-master JPL (MM-JPL). On the base station, we lever-




Figure 2. System architecture




Figure 3. Main mission state machine described with the BPMN specification

age web technologies to provide intuitive monitoring and controlling interface for the human supervisor.

Mission-level Autonomy

The Traceable Robotic Activity Composer and Executive (TRACE) [15] is the mission executive running on each robot that executes the mission specified in a Business Process Model and Notation (BPMN) file as shown in Figure 3. TRACE controls the progression of the mission in a manner similar to a finite state machine and also accepts operator commands to end the current mission or load a new mission. Each block in the BPMN file corresponds to a behavior, of which there are currently two types: mobility behaviors (such as transiting to a frontier) and monitoring behaviors (such as monitoring the tilt of the vehicle). The BPMN notation allows the mission designer to have multiple behaviors executing within the mission simultaneously through the use of Parallel Gateways. In addition, an independent Mission Watchdog

is responsible for overriding the current mission at certain events in order to ensure that the supervisor receives new data from the robot before the end of the competition.

The following behaviors were used during the competition: 1) Transit to Frontier, for moving to a specified frontier point using IRM nodes as intermediate waypoints, 2) Push Frontier Forward, for exploring new areas in the forward direction, 3) Return to Base, for returning to the base IRM node, 4) Monitor Tilt, for monitoring the pitch and roll of the robot, 5) Monitor Heartbeat, for monitoring when communications to the base station have been restored, and 6) Move Backwards, a contingency behavior that is activated when the default waypoint following fails to bring the robot back to base, perhaps due to localization drift. The Monitor Tilt behavior runs throughout the mission and immediately ends the mission when it has detected that the robot has either rolled or pitched past a certain threshold. The Monitor Heartbeat behavior only runs simultaneously with the Return to Base behavior and stops the robot when the human supervisor re-establishes communications with the robot. The combination of these behaviors results in a Return to Communications behavior and enables the robot to remain far inside the environment. This saves the operator time as the robot does not have to come back all the way to base before being commanded another frontier to explore.

Finally, the Mission Watchdog was triggered at the following events: 1) Communications Timeout, to force the robot to return to communications range after a specified time spent without communications with the base, 2) Game-Over Timeout, to force the robot to return to communications range before the end of operations, and 3) Idle Timeout, to force the robot to attempt to return to base if it has not been moving for a while (perhaps due to a temporary blockage).




Figure 4. Communications architecture

Global Path Planning and Exploration

The global planner relies on a graph that captures the connectivity of the free space. We refer to this graph as IRM (information roadmap) as its edges and nodes are enriched by various environmental and robot-related features. IRM used in the first phase of the DARPA competition is a simplified version of the information roadmap introduced in our previous work [18], [19], [20].

The global planner provides services to maintain and query the IRM. For example, behaviors can query the global planner for the shortest path to frontiers and the base station. Additionally, the planner is responsible for ensuring that the locations of the IRM nodes remains consistent after a pose-graph relaxation from the global localization module. Each IRM node is anchored to a pose-graph node on creation; when the pose-graph node updates its pose estimate after a loop closure, the planner applies the same transformation to all IRM nodes that are anchored to it. When a loop-closure occurs (either triggered manually or autonomously), the planner will also update the edges between IRM nodes to reflect the new topology of the pose graph.

The human supervisor initially constructs the IRM with a base node and a frontier node. This roadmap is then sent to the global planner running on each robot. Next, the Transit to Frontier behavior will query the global planner for a path to the assigned frontier node. Once the robot has reached the frontier, it will proceed to start the exploration behavior. The following tasks simultaneously modify the IRM at this time: 1) breadcrumb dropping, 2) frontier sampling, and 3) frontier pruning. Breadcrumb dropping creates new IRM nodes along the path traversed by the robot and allows it to find a simple path back to base. Frontier sampling periodically places frontier nodes in a semi-circle in front of the robot which are collision-checked against a traversability cost map. Finally, frontier pruning removes existing frontier nodes that have been considered ‘explored’—i.e., if the frontier node has been covered by the robot’s trajectory, expanded by a dynamically changing radius. The dynamically changing radius prevents too many frontiers from being created in wider corridors while still preserving frontiers to smaller openings.

A terminal IRM node is created when the robot detects a dead-end (detected by the lack of feasible paths to a goal in front). At this point, the Push Frontier Forward behavior will select the next closest frontier to transit to and explore next.



Figure 5. Two communication systems tested at Beckley Exhibition Mine with Husky robots. Left: tether, Right: communication dropper

Communications

The robots communicate with the base by means of a layer-2 mobile ad-hoc mesh network (MANET). For the Tunnel Circuit, commercial off-the-shelf MANET radios from Silvus Technologies (SC4240E-235-BB) and Persistent Systems (MPU5) were fielded on each robot. Although the radios offer a range of frequencies of operation, they were operated within federal (NASA) and Industrial, Scientific and Medical (ISM) bands for which a short term use Special Temporary Authorization (STA) license was obtained.

Since the mine environment is by nature non-line-of-sight, communications nodes are dropped by the robots, incrementally building a backbone network for reach back to the base station (see Figure 4 for a conceptual diagram). Communication nodes are created by placing a droppable radio on its side and adding 90 degree adapters for the antennas. The nodes are deployed from two different dropping mechanisms, one of which is shown in Figure 5. This one uses a stack of radios in a track with pairs of solenoids controlling the node release. The other comm node dropper uses a motor to move a tray of comm nodes closer to the ground and servos to control individual radio release.

The node drop decision was based on SNR thresholding. The SNR may fall to the level of this threshold simply because of distance traveled or bends in the tunnel. In narrower mines with rough walls, such as a gold mine, the signal reaches less far, but the scattering helps address communications when there is wander in the passage way and also provides some communications around corners. In smoother passageways, the signal is able to carry further due to the shallow incidence angles with the electrically flat walls, but offers less ability to scatter down any side passage ways. To achieve the largest possible distance per node, high power settings and Multiple-in-multiple-out (MIMO) were employed.

A 2-wire 300m Digital Subscriber Line (DSL) tether was also deployed from a robot during field tests. As seen in Figure 5, the tether deployment system consisted of a 300m spool of 2 conductor wire are mounted on a Clearpath Husky robot. The spool was connected to a DSL modem in the Husky via a slip ring mounted on the side of the spool. A spool brake provided some friction to stop the spool from moving if the robot stopped.

The multi-master JPL (MM-JPL) software was developed to efficiently exchange messages between the base and the robots over the network. It was implemented as ROS nodes running on each robot and on the base station. The inter-robot communications over the MANET was implemented via a per ROS topic Selective Repeat ARQ (automatic repeat-request) over the User Datagram Protocol (UDP) with rate metering. Vehicle status, artifact reports and mapping information were sent back to the base, and commanding of the robots was performed from the base. Although MM-JPL was configured to convey these messages, MM-FKIE was also available for increased flexibility and for some other message types deemed less mission critical.

Operation Tools

Human-system interaction is an essential process in any mission operation including the DARPA SubT Challenge. We developed a set of visualization and commanding interfaces that are tailored not only to support multi-robot coordination, but also specifically for executing underground operations. We describe the main design features and architecture of the set of user interface (UI) tools developed for the operations team during field tests and the Tunnel Circuit.

The main goal of the UI tools is to provide the supervisor/operator (and his/her supporting team) the capability of 1) monitoring the state and health of robots and the entire system, 2) monitor the progress of the mission, 3) dispatching commands to the team of robots, and 4) managing critical target data such as the detailed 3D map of the subsurface environment. In the context of the SubT Challenge, artifact location and type is the most relevant and time critical information.

In addition to designing the tools towards meeting this goal, we also shaped them according to operational requirements defined by the Challenge. There is a hard constraint on the size of the team that can actively operate (interact with) the system: only *one* operator. The single-operator setting is quite challenging when there are several moving parts to manage during a mission (e.g., high volume of data and/or info to manage, high cognitive load for operating robots simultaneously). Our current set of tools addresses some of these unique challenges. As part of requirements, pit crews are allowed to support the single supervisor with robot set up and limited situational awareness; however, no interaction with the operation tools is allowed. The design of tools reflects that separation of roles (supervisor and pit crew). Below we describe the current design to meet the aforementioned constraints and requirements.

Framework—We leverage ROS tools and common web-based technologies for user interfacing and data storage. To interact with robot’s ROS software, we developed a UI Manager node that was responsible for receiving robot status, map data, network measurements and artifact status information from robots and storing them in a MongoDB² database. The UI Manager is also responsible for translating high-level commands from the supervisor to the robots, also through ROS. The situational awareness and commanding tools are implemented using a series of React-based³ web-pages, with a Node.js⁴ web-server as the back-end. These tools both support human interaction with the entire system, and allow for data manipulation, such as the alteration of properties.

²<https://www.mongodb.com>

³<https://reactjs.org>

⁴<https://nodejs.org>




Figure 6. A snapshot of operation tools

This server was also designed to manage the interaction between the base station and the DARPA scoring system when submitting new artifacts. The submission status are also stored in the database.

Views—The collected information is presented to the supervisor through web-based views and ROS RViz-based 3D views. In what follows, we describe each of the main views designed for operations.

3D Spatial View: Situation awareness in subsurface environments demands a rich set of information available to the human supervisor in real time. One of the central UI to meet that demand is the 3D Spatial View (Figure 6 left side, RViz with plugins) where the operator can observe the most up to date 3D point cloud of the environments, the current pose graph (blue nodes) and IRM graph (orange nodes) along with frontier node candidates (yellow cubes) branched out from IRM breadcrumb nodes, and the current position of the robots. Through that interface, the supervisor is also able to send commands to robots to navigate to specific nodes in the IRM graph or adjust the 3D map to fix potential drifts and inconsistencies in localization.

Mission and Robot State View: To accommodate the massive volume of status data arriving to the base, we developed the Mission/Robot State View (Figure 6 right side) to serve as the main web-based UI for monitoring each robot’s health and mission state, and for commanding the team of robots. Figure 6 shows three colored sections on the right view, one per robot. Each section shows the robot’s health status as a set of icons with the major/critical sensing information (e.g., battery level, network connectivity) and a list of errors and warnings produced by the robot’s ROS components and modules (navigation system, communications, database, etc). Each robot’s section also provides a set of commands to control the robot, such as uploading, starting and stopping a BPMN mission to be executed (in which the status of execution is also monitored by showing active behaviors in green), as well as dropping a new communication node. Moreover, such robot-centric view also provides on-demand image feeds from onboard cameras pointing multiple directions. During operations the Mission/Robots State View is the most used UI by the supervisor.

Network and Data View: The Network and Data View pro-

vides network situation awareness to the operator to support communication/network diagnosis and comm node dropping based on SNR and data rate readings. The view also presents the real-time network topology, where robot pose is used to arrange the visual nodes in the graphical representation of the network, and SNR values are used to create edges between those graphical nodes.

Artifact View: The Artifact View is a major tool for operations during the SubT Challenge runs. We carefully designed this view to support the processes around artifact data management, including information arrival to the base station (through UI Manager and database), inspection and analysis by the supervisor, submission to the DARPA scoring server, and status maintenance. Artifacts are shown initially through the list of detected artifacts, along with their thumbnail picture, classified type, position, and confidence level. The supervisor can inspect that information by clicking on one of the items in the list and augmenting that single artifact information to better inspect the data and the picture taken by the robot. We also developed an elaborated feature for searching and sorting discovered artifacts to provide a fast way to browse detected artifacts using their attributes.

If the artifact data is approved by the human supervisor, the UI assists with the process of submitting the data to DARPA server for scoring. The response from the DARPA server is stored in the database and shown to the user (e.g., scored or unscored). The status and search mechanisms help the supervisor to navigate through both artifacts that need to be submitted and those that might require inspection due to rejection (in this case the classification or the localization of the item might be off).

Dashboard View: We developed a Dashboard view to summarize all the essential status data across all the aforementioned views. This particular view is also used to manage mission duration and the number of artifacts still to be discovered. This summarization becomes quite important in time critical mission management so that the operation team can quickly evaluate the state of the system and take actions.

Pit Crew Views: As mentioned above, the interaction of the pit crew with the system is limited to visualization only (no keyboard access) during the runs. Therefore, we designed a set of views to provide key information to particular pit crew members (such as communication/networking, hardware, mapping and localization, object detection and autonomy) but restricted the visualization to only show data which is allowed by the DARPA Subterranean Challenge rules (e.g., no map or artifact information can be shared with the pit crew visually through these views).

4. RESULTS

The presented approach was deployed onto a team of ground and aerial vehicles operated by a single human supervisor. We conducted a series of field demonstrations at 9 different sites over the past year, including the Edgar Experimental Mine during the SubT Integration Exercise (STIX) event and the National Institute for Occupational Safety and Health (NIOSH) mine complexes near Pittsburgh, PA for the Tunnel Circuit. We experimented with different robot configurations in each test as our technologies matured over time. As of the Tunnel Circuit, based on available autonomy capabilities and human’s cognitive load, we converged to use two ground vehicles as Vanguard, and 1–4 ground/aerial vehicles as Support



Figure 7. Four Husky robots standing by at the gate of Safety Research mine during the Tunnel Circuit event.




Figure 8. IRM overlaid with map from tunnel circuit.

depending on the environmental complexity. Figure 7 shows our robot team waiting for the run start at the Safety Research mine during the Tunnel Circuit. Two vanguard robots are followed by two support robots carrying communication node droppers. In this section, we summarize the results and lessons learned throughout the field test campaigns.

Autonomy

Mission Design and Execution—The main mission used during the competition for each robot is shown in Figure 3. This mission is common to all robots, except that Support robots receives more interventions from the supervisor for tasks such as communication backbone maintenance. Since all robots were running the same mission, operator intervention was required at times to assign robots to different frontiers to ensure more complete coverage of the environment. This could be alleviated by having an autonomous coverage coordinator, or assigning different frontier selection policies.

The Return to Communications behavior was particularly effective in the field. All robots successfully returned within the communications range and transferred map data before the end of each run; the longest amount of time any robot spent outside of communications range was 800 seconds.

IRM construction—Figure 8 illustrates the IRM graph and corresponding map constructed by one of the robots during the Tunnel Circuit. The breadcrumb nodes are shown as orange spheres, and the frontiers are shown as yellow cubes.




Figure 9. Maps from Tunnel Circuit (Left: Experimental Mine, Right: Safety Research Mine). The communication node drop locations are marked by crosses.

In this region, the robot was fully autonomous without any communication to the base. The current frontier selection policy let the robot choose the frontiers in front, while dropping frontier nodes in the side passages for later exploration. If the robot detects a dead-end (marked by red cubes), it continues exploration by selecting the closest unexplored frontier.

Triggered by a loop closure event by the global localization module, the IRM breadcrumb nodes were rewired to capture the loop structure. In this area in Figure 8, as a result of the IRM rewiring, the robot was able to return to base/communications with much shorter distance, saving 300 m extra drive at minimum. However, we also observed in the field tests that an improper loop closure might cause IRM rewiring failure which places an IRM node inside a wall; as a result, the robot had difficulty coming back to the base. The robustness to loop closure failure is listed as our future work.

While the IRM construction process was mostly done autonomously, there were some cases human assistance was needed. One of such examples is the missing frontier placement in passages that are too narrow or located outside of the frontier sampling ranges. If that happens, our approach is to let the supervisor add a new frontier node through the UI and send the modified IRM to the robot to guide the exploration behavior. This strategy worked well in most of tunnel environments, as long as we can maintain a good communication backbone to send back the volumetric map for human inspection.

Communications

Tethered Communications—Wired link through the DSL tether gave the stable communications between the robot and the base station. We were even able to tele-operate the tethered robot from the base station while streaming camera images at a high rate. The most challenging part of the tethered system was the deployment. In our experiment, the initial deployment of the tether worked extremely well. However, the combination of mud in our deployment system and jerky robot motion in a stop-and-go manner was found to cause the tether to lasso the deployment system. A more sophisticated deployment system could solve these issues, but was left for future work.

Wireless Communications—Both Silvus Technologies and Persistent Systems radios performed well in the smooth walled tunnels of the NIOSH mines. Communication over

the maximum line-of-sight distance of 400 m was achieved. Unlike some mines where we observed signal bouncing around the corners, droppable comm nodes were necessary to allow communication into side passages as indicated by cross marks on Figure 9.

During the field test campaigns, we occasionally observed instability of network (e.g., large latency, high packet loss rate). Our inspection concluded that this instability was caused by excessive bandwidth usage. We identified two data sources that contribute to the excessive data volume. Firstly, there were unintentional data flows due to misconfiguration in message syncing. Monitoring any data flow between robots could have helped catching such unexpected traffic earlier. Secondly, the increasing map data volume of the large-scale environment eventually exceeds the threshold that the network can handle; as a result the performance was drastically degraded. We have to be strict about what needs to be sent to the network and how.

Human-system Interaction

As our technologies evolved, the frequency of human intervention became lower. In the Tunnel Circuit, the human interaction is mostly related to high-level mission commanding or guiding exploration direction. However, the supervisor occasionally sent lower-level commands (e.g., short-range mobility command) to obtain better performance under the time and communication constraints.

We measured the frequency of human interaction during the third run at the Safety Research mine. We deployed two Vanguards and one Support for this run. The supervisor sent 16 missions in total, including 4 comm node dropping missions. In the current implementation, each comm node drop requires the position alignment by the supervisor to achieve better signal performance. The position alignment took 1–3 minutes depending on the distance and available bandwidth. The supervisor edited the IRM four times, spending ~40 seconds per each. The primary reasons for editing the IRM were 1) to add missing frontiers and 2) to force exploration direction for better coverage.

5. CONCLUSIONS

We presented our approach for supervised multi-agent autonomy to explore and map subsurface environments. The robot team was able to explore unknown communicationally

degraded environment with minimal but effective help of the human supervisor. We evaluated the performance of our system in the context of DARPA SubT Challenge. The robot team was able to explore kilometers of regions while constructing an accurate map with only a few meter error.

ACKNOWLEDGMENTS

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. This work was partially funded by the Defense Advanced Research Projects Agency (DARPA). Government Sponsorship acknowledged.

REFERENCES

- [1] P. J. Boston, R. D. Frederick, S. M. Welch, J. Werker, T. R. Meyer, B. Sprungman, V. Hildreth-Werker, S. L. Thompson, and D. L. Murphy, "Human utilization of subsurface extraterrestrial environments." *Gravitational and Space Biology Bulletin, Publication of the American Society for Gravitational and Space Biology*, vol. 16, no. 2, pp. 121–31, 2003.
- [2] G. E. Cushing, "Candidate cave entrances on Mars," *Journal of Cave and Karst Studies*, vol. 74, no. 1, pp. 33–47, 2012.
- [3] F. Hörz, "Lava tubes: Potential shelters for habitats," *Lunar Bases and Space Activities of the 21st Century*, pp. 405–412, 1985.
- [4] P. J. Boston, M. N. Spilde, D. E. Northup, L. A. Melim, D. S. Soroka, L. G. Kleina, K. H. Lavoie, L. D. Hose, L. M. Mallory, C. N. Dahm, L. J. Crossey, and R. T. Schelble, "Cave biosignature suites: Microbes, minerals, and Mars," *Astrobiology*, vol. 1, no. 1, pp. 25–55, 2001.
- [5] R. J. Léveillé and S. Datta, "Lava tubes and basaltic caves as astrobiological targets on Earth and Mars: A review," *Planetary and Space Science*, vol. 58, no. 4, pp. 592–598, 2010.
- [6] DARPA Subterranean Challenge. [Online]. Available: <https://www.subtchallenge.com>
- [7] DARPA Subterranean Challenge Team CoSTAR. [Online]. Available: <https://costar.jpl.nasa.gov>
- [8] W. Walsh and J. Gao, "Communications in a cave environment," in *IEEE Aerospace Conference*, 2018, pp. 1–8.
- [9] A. Husain, H. Jones, B. Kannan, U. Wong, T. Pimentel, S. Tang, S. Daftry, S. Huber, and W. L. Whittaker, "Mapping planetary caves with an autonomous, heterogeneous robot team," in *IEEE Aerospace Conference*, 2013, pp. 1–13.
- [10] T. Vaquero, M. Troesch, and S. Chien, "An approach for autonomous multi-rover collaboration for Mars cave exploration : Preliminary results," in *International Symposium on Artificial Intelligence and Robotics & Automation in Space*, 2018.
- [11] T. Vaquero, M. Troesch, M. S. Net, J. Gao, and S. Chien, "Energy-aware data routing for disruption tolerant networks in planetary cave exploration," in *International Workshop on Planning and Scheduling for Space*, 2019.
- [12] A. Santamaria-navarro, R. Thakker, D. D. Fan, B. Morrell, and A. Agha-mohammadi, "Towards resilient autonomous navigation of drones," in *International Symposium on Robotics Research*, 2019.
- [13] K. Ebadi, Y. Chang, A. Stephens, A. Hatteland, E. Heiden, A. Thakur, B. Morrell, L. Carlone, and A. Agha-mohammadi, "LAMP: Large-scale autonomous mapping and positioning for exploration of perceptually-degraded subterranean environments," in *IEEE International Conference on Robotics and Automation*, 2020 (under review).
- [14] E. Terry, X. Lei, B. Morrell, S. Daftry, and A. Agha-mohammadi, "Artifact detection and localization in perceptually-degraded subterranean environments," in *IEEE International Conference on Robotics and Automation*, 2020 (under review).
- [15] J.-P. de la Croix, G. Lim, J. Vander Hook, A. Rahmani, G. Droge, A. Xydes, and C. Scrapper Jr, "Mission modeling, planning, and execution module for teams of unmanned vehicles," in *Unmanned Systems Technology XIX*, vol. 10195. International Society for Optics and Photonics, 2017, p. 101950J.
- [16] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Ng, "ROS: An open-source robot operating system," in *IEEE International Conference on Robotics and Automation, Workshop on Open Source Software*, 2009.
- [17] A. Tiderko, F. Hoeller, and T. Röhling, "The ROS multimaster extension for simplified deployment of multi-robot systems," *Robot Operating System (ROS), Studies in Computational Intelligence*, vol. 625, pp. 629–650, 2016.
- [18] A. Agha-mohammadi, S. Chakravorty, and N. Amato, "FIRM: Sampling-based feedback motion planning under motion uncertainty and imperfect measurements," *International Journal of Robotics Research (IJRR)*, vol. 33, no. 2, pp. 268–304, 2014.
- [19] A. Agha-mohammadi, S. Agarwal, S. Kim, S. Chakravorty, and N. M. Amato, "SLAP: Simultaneous localization and planning under uncertainty via dynamic replanning in belief space," *IEEE Transactions on Robotics*, vol. 34, no. 5, pp. 1195–1214, 2018.
- [20] S. Kim, R. Thakker, and A. Agha-Mohammadi, "Bi-directional value learning for risk-aware planning under uncertainty," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2493–2500, 2019.

BIOGRAPHY



Kyohei Otsu is a Robotics Technologist at NASA Jet Propulsion Laboratory, California Institute of Technology. He received his B.S., M.S., and Ph.D. degrees in Electrical Engineering from the University of Tokyo. He is interested in autonomous navigation technologies and robotic system architecture.



Scott Tepsuporn is a Robotics Electrical Engineer within the Mobility and Robotic Systems section at the Jet Propulsion Laboratory where he is involved in the design and implementation of motion planning algorithms for unmanned surface vessels (USV) and autonomous under-water vehicles (AUV). He has also worked on a distributed communications framework for collaborative multi-agent USV projects. He received his B.S. degrees from the University of Virginia in Computer Science and Electrical Engineering, where he worked on high-speed multi-domain data transfers.



Rohan Thakker is a Robotics Technologist at Nasa's Jet Propulsion Laboratory, California Institute of Technology. He received his B.S. degree from National Institute of Technology, Nagpur, India and M.S. from Carnegie Mellon University, Pittsburgh. Rohan is interested in developing robust autonomous navigation technologies by tightly coupling perception/estimation, motion planning and controls.



Tiago Vaquero is a Technical Group Leader in the Artificial Intelligence Group, Planning and Execution Section, of the Jet Propulsion Laboratory, California Institute of Technology, where he works on multi-agent coordination methods for multi-rover cave exploration and surface site characterization. Tiago holds a B.Sc., M.Sc., and Ph.D. in Mechatronics Engineering from the University of Sao Paulo, Brazil. Tiago previously held a MIT research scientist position and a joint Caltech/MIT research position where he worked on Resilient Spacecraft Systems and Risk Sensitive Planning/Scheduling algorithms. At MIT, Tiago also worked on Risk Sensitive Planners and Executives for Autonomous Underwater Vehicles and Autonomous Cars. Tiago also held a research position at the University of Toronto where he worked on Robotics for Healthcare and Multi-Robot Coordination.



Jeffrey A. Edlund is a Robotics Systems Engineer at NASA's Jet Propulsion Laboratory. He received his B.S. in Physics and Ph.D. in Computation and Neural Systems from the California Institute of Technology.



William Walsh holds Bachelor degrees in Physics and Math from the University of Minnesota and a Masters of Science in Physics from UCLA. His area of work has been in system and network architectures for commercial and government space based communications systems. At NASA he has worked on the Constellation Program for human space flight and on networked services for high rate optical communications (LCRD program). Prior to JPL, he

as been a Corporate Engineer at Linquest Corporation, and an associate technical fellow at Boeing.



Gregory Miles is a Telecommunications Engineer at NASA's Jet Propulsion Laboratory (JPL) in the Communications Architecture and Research Section. At JPL he works on software development for ground data systems, embedded systems and Delay Tolerant Networks (DTN). He received his B.S. in Computer Science from California State University Los Angeles.



Tristan Heywood was a Robotics Technologist Intern at NASA Jet Propulsion Laboratory, California Institute of Technology, for most of 2019. He is currently completing a B.S. in Computer Science and a B.E. in Mechatronic Engineering. Tristan worked on the development of the operation tools.



Michael T. Wolf is a Principal Robotics Technologist and the supervisor of the Maritime & Multi-Agent Autonomy group at the NASA Jet Propulsion Laboratory (JPL). He earned his M.S. and Ph.D. in Mechanical Engineering (Robotics) at the California Institute of Technology and his B.S. in Mechanical Engineering at Stanford University. His research interests include multi-agent systems, autonomy software architectures, maritime robotics, and vehicle autonomy.



Ali-akbar Agha-mohammadi is a Robotics Research Technologist with NASA's Jet Propulsion Laboratory (JPL), Caltech. Previously, he held a position of an Autonomy Research Engineer with Qualcomm Research and a Postdoctoral Researcher with the Laboratory for Information and Decision Systems at MIT. He received his Ph.D. from Texas A&M University. His research interests include robotic autonomy, mobility and perception, stochastic control systems, and filtering theory. Dr. Agha leads several projects at JPL on robotic autonomy, including JPL's CoSTAR team participating at the DARPA SubTERRANEAN Challenge. In 2018, Dr. Agha selected as a NASA NIAC Fellow.